arXiv:1401.0074v1 [gr-qc] 31 Dec 2013

Perturbative stability of the approximate Killing field eigenvalue problem

Christopher Beetle and Shawn Wilder

Department of Physics, Florida Atlantic University, 777 Glades Road, Boca Raton, Florida 33431

Abstract. An approximate Killing field may be defined on a compact, Riemannian geometry by solving an eigenvalue problem for a certain elliptic operator. This paper studies the effect of small perturbations in the Riemannian metric on the resulting vector field. It shows that small metric perturbations, as measured using a Sobolev-type supremum norm on the space of Riemannian geometries on a fixed manifold, yield small perturbations in the approximate Killing field, as measured using a Hilbert-type square integral norm. It also discusses applications to the problem of computing the spin of a generic black hole in general relativity.

1. Introduction

The gravitational field in general relativity does not admit a local stress-energy tensor. Physically observable quantities like energy and angular momentum therefore are defined only quasi-locally, *i.e.*, as integrals over 2-dimensional surfaces in 4-dimensional spacetime. Moreover, these quasi-local integrals are completely unambiguous only when spacetime admits a continuous symmetry generated by a Killing vector field. Once justified in the presence of symmetry, however, the integrals can also be calculated more generally, provided that one has (a) a means of choosing a particular 2-dimensional surface over which to integrate and (b) a means of picking a particular vector field on that surface to play the role previously played by the Killing field. The latter issue is particularly troublesome because a typical manifold has no exact symmetries whatsoever. As a result, several schemes [1–5] have appeared in the literature in recent years to define *approximate* Killing vector fields on manifolds that have no real symmetries. This paper is concerned with one such scheme, wherein the approximate Killing field arises by solving a certain eigenvalue problem on a given Riemannian geometry [4–6]. More precisely, here we investigate the *stability* of this scheme under small perturbations of the geometry. The question is important because one would expect that any physically reasonable definition of, say, the angular momentum of the gravitational field in a region of spacetime should vary continuously under a continuous deformation of that region, or of the geometry within. We show below that the approximate Killing vector eigenfield depends continuously, in a precise sense, on the

geometry used to define it. This implies that any quasi-local integral based on this approximate Killing field varies continuously as well.

The detailed calculations that follow have an unavoidably mathematical tenor. Nonetheless, they are motivated primarily by practical concerns arising in physical applications of the approximate Killing field eigenvalue problem. The main such application is to the problem of calculating the spin of a black hole [3–5]. Let us review this application since it is the essential conceptual backdrop for the present results.

The angular momentum of a globally stationary and axisymmetric black-hole spacetime is given by the Komar formula [7, 8]

$$J_{\text{Komar}}[S,\varphi] := \frac{1}{16\pi G}\oint_S \star \mathrm{d}\varphi. \tag{1}$$

The right side here is the integral of the (spacetime) Hodge dual of the exterior derivative of $\varphi_a := g_{ab}\,\varphi^b$ over a 2-sphere S enclosing the black hole. The vector field used to calculate the Komar angular momentum is a *global* axial Killing field on spacetime. For vacuum spacetimes, $J_{\text{Komar}}[S,\varphi]$ is independent of the 2-sphere S, and coincides with the total angular momentum of spacetime measured by inertial observers at infinity. It is natural to identify this angular momentum with the spin of the black hole in this case because the assumed symmetries preclude any gravitational radiation in the intervening space that might contribute additional angular momentum.

The integrals (1) certainly exist for non-symmetric spacetimes. But of course they depend one's choices of both S and φ^a. These ambiguities can be resolved in certain circumstances, such as if (i) one can identify the sphere S with a cross-section either of the boundary of a 4-dimensional region $\mathcal{R}$ of spacetime or of a black hole's horizon, and (ii) the *intrinsic* geometry on S admits an axial Killing field ξ^b. Various *quasi-local angular momentum* formulae have been proposed for such circumstances. Naturally, these must reduce to the Komar angular momentum (1) if there exists a global axial symmetry φ^a such that $\xi^b = \varphi^b\big|_S$ on the 2-sphere in question. (The restriction $\varphi^b\big|_S$ must then be tangent to S.) If such an extension does not exist, however, then the integrand for a quasi-local angular momentum may differ from that of (1).

Consider, for example, the *Brown–York angular momentum*

$$J_{\text{BY}}[S,\xi] := -\oint_S {}^{\mathcal{T}}\star\,(\xi \lrcorner \tau). \tag{2}$$

The 2-sphere S here is a cross-section of a "tubular," time-like, outer boundary $\mathcal{T}$ of a 4-dimensional region $\mathcal{R}$ of spacetime. Meanwhile, ${}^{\mathcal{T}}\star$ denotes the 3-dimensional Hodge dual within $\mathcal{T}$, and $(\xi \lrcorner \tau)_b := \xi^a\,\tau_{ab}$ involves the extrinsic curvature τ_{ab} of $\mathcal{T}$ within spacetime. Another example is the *generalized angular momentum*

$$J_{\text{Gen}}[S,\xi] := \frac{1}{8\pi G}\oint_S {}^{\Sigma}\star\,(\xi \lrcorner K) \tag{3}$$

motivated by the study of dynamical horizons [9, 10]. The 2-sphere S in this case lies in a given partial Cauchy slice Σ, and is such that the expansion of one of the two

null geodesic congruences emanating from it vanishes. Meanwhile, ${}^{\Sigma}\star$ denotes the 3-dimensional Hodge dual within Σ, and K_{ab} denotes the extrinsic curvature of Σ as it is embedded in spacetime.

The common feature of (2) and (3) is that the sphere S is fixed geometrically by features of the surrounding spacetime. Both formulae are then on their surest footing when the intrinsic geometry on S admits an axial Killing field ξ^a. To continue interpreting (2) or (3) as *the* angular momentum of a black hole when S is not axially symmetric, one must specify how to construct a similarly geometrically preferred vector field ξ^a from an arbitrary 2-sphere metric. Moreover, this vector field must reduce to the axial Killing field whenever one exists. A natural way to do this is to seek a best approximation, in some prescribed sense, to a Killing field on S. Several schemes of exactly this type have been proposed [1–6].

In principle, prescribing the two structures (a) and (b) gives a precise analytic formula for a black hole's angular momentum. But subtleties can arise in practical applications, such as in numerical relativity. In particular, the preferred 2-sphere S can be expensive to calculate numerically [11]. Thus, the sphere S one actually uses in practice may be displaced slightly, and have a slightly different intrinsic geometry, from the ideal sphere on which the analytic formula is based. If, in addition, the choice of approximate Killing field depends on some *ad hoc* choices, such as in [1, 2], then one must also worry about how those choices affect the angular momentum. Although the approximate Killing field derived from the eigenvalue problem described below does not suffer from the latter ambiguity, it is not immediately clear how important it is to locate the desired S exactly in spacetime.

This paper shows explicitly that small errors in the ideal intrinsic geometry on S produce only small errors in the approximate Killing field determined by the eigenvalue problem described below. Importantly, the quasi-local angular momentum formulae (2) and (3) both lead to expressions like

$$J[S] = \oint_S \xi^a[g]\, j^b\, g_{ab}\, \epsilon. \tag{4}$$

This is an L^2 inner product of an approximate Killing field $\xi^a[g]$, constructed (non-locally) from the intrinsic geometry g_{ab} of S, with a vector field j^b that depends (locally) on both the intrinsic and extrinsic geometry of S. The change of any such integral under perturbations of the fields involved can be bounded using the Schwarz inequality. The question thus arises whether the L^2 norm of the perturbation in $\xi^a[g]$ can be bounded in terms of perturbations of the intrinsic metric itself, and therefore in terms of potential errors locating S exactly within spacetime. We show below that it can.

Although our primary motivation is the application to the 2-dimensional geometry of the horizon of a black hole, the results below actually apply to perturbations of a Riemannian geometry of arbitrary dimension. The topology of the manifold, however, must remain compact and without boundary. One could presumably relax the latter restriction by imposing appropriate boundary conditions, but we will not explore this question in detail.

The outline of the paper is as follows. Section 2 reviews the eigenvalue problem from [4–6] whose solution yields the approximate Killing field. Section 3 describes a novel scheme to fix the (global) normalization of that vector eigenfield. Section 4 discusses perturbations of the eigenvalue problem, and shows that perturbations in the eigenfields can be bounded if the perturbation of the underlying operator can be bounded appropriately. Section 5 shows that the perturbed operator can indeed be bounded in this way. We conclude in Section 6 with some comments and discussion.

2. The Killing Laplacian and approximate Killing fields

Let $\mathcal{M}$ be a compact Riemannian manifold with metric g_{ab}, and let u^a denote an arbitrary vector field on that manifold. Define the action integral

$$\mathcal{S}[u](\kappa) := \frac{1}{2}\int_{\mathcal{M}} \mathscr{L}_u g_{ab}\, \mathscr{L}_u g_{cd}\, g^{ac}\, g^{bd}\, \epsilon + \kappa\left(1 - \int_{\mathcal{M}} u^a\, u^b\, g_{ab}\, \epsilon\right), \tag{5}$$

where ϵ denotes the volume form on $\mathcal{M}$ derived from g_{ab}. Minimizing this integral obviously minimizes the Lie derivative of the metric in an averaged sense over the whole of $\mathcal{M}$, subject to a constraint that the natural L^2 norm of the vector field u^a on $\mathcal{M}$ remains fixed. This rules out the limit $u^a \to 0$, where clearly (5) always has a trivial minimum. The Euler–Lagrange equation for (5) is

$$\Delta_{\mathrm{K}}\, u^d := -2\, \delta^{(d}_c\, g^{a)b}\, \nabla_a\, \nabla_b\, u^c = \kappa_u\, u^d. \tag{6}$$

Its solutions are vector eigenfields u^a for the *Killing Laplacian* operator Δ_{K} with corresponding eigenvalues κ_u. Historically, it appears that this operator was first introduced to the relativity literature by Matzner [6], and independently to the differential geometry literature by Bochner and Yano [12]. More recently, Owen and collaborators [4] have reintroduced the idea in specific application to the angular momentum of black holes.

The Killing Laplacian Δ_{K} has several important characteristics that dictate the nature of its eigenvalue problem. First, it is positive and Hermitian on the Hilbert space $L^2(\mathcal{M}^\bullet, g)$ of vector fields‡ on $\mathcal{M}$ that are square integrable with respect to the natural measure defined by g_{ab}. The eigenvalues κ_u are therefore real and positive, and the corresponding eigenfields u^a are real and orthogonal in the inner product

$$\langle u, v\rangle_{L^2(\mathcal{M}^\bullet, g)} := \oint_{\mathcal{M}} \bar{u}^a\, v^b\, g_{ab}\, \epsilon. \tag{7}$$

Second, the Killing Laplacian is (strongly) elliptic [13]. Such operators have a *complete* basis of *normalizable* eigenfields u^a with *discrete* eigenvalues κ_u that are *unbounded* above.

‡ The superscript on $\mathcal{M}$ indicates the tensorial structure of the fields in the Hilbert space $L^2(\mathcal{M}^\bullet, g)$, and we include g explicitly to emphasize that the inner product varies with the metric. We will drop the latter notation shortly, but keep the former throughout because several Hilbert spaces of different types of tensor fields appear in our calculations.

We define the *approximate Killing vector field* for an arbitrary Riemannian geometry $(\mathcal{M}, g_{ab})$ to be simply the eigenfield u_0^a of the Killing Laplacian with the smallest eigenvalue $\kappa_0 \geq 0$. If g_{ab} admits an actual Killing field ξ^a, then u_0^a coincides with it and has eigenvalue $\kappa_0 = 0$. But of course the definition is valid generally. It also may happen that the lowest eigenvalue κ_0 of Δ_{K} is degenerate, and u_0^a therefore may not be unique. (This happens, for example, for a round sphere, in which case the lowest eigenvalue is $\kappa_0 = 0$ and the corresponding eigenspace is 3-dimensional. We do not have a similarly concrete example of a geometry with $\kappa_0 > 0$ for which the corresponding eigenspace is multi-dimensional, but there is no obvious reason to suppose that such geometries do not exist.) We will not consider this possibility explicitly below, however, because the generic results for non-degenerate κ_0 that we do present generalize straightforwardly to this exceptional case.

3. Normalizing the approximate Killing field on a sphere

One of course has to normalize the approximate Killing field u_0^a appropriately in order to calculate a definite spin for a black hole. As for any eigenfield, the only freedom is to scale it by a constant over all of $\mathcal{M}$. But it is not immediately clear how to choose that constant scaling. The normalization for which the constraint term vanishes in the action principle (5) just fixes the Hilbert norm of u^a in $L^2(\mathcal{M}^\bullet, g)$. This is almost certainly not the correct normalization if we are trying to model an axial Killing field on a topological 2-sphere. (It would be a fine normalization for a translational Killing field on a torus, say, but this is not the case that interests us most.) The question therefore arises how to normalize u_0^a when $\mathcal{M} \sim S^2$ such that it reduces to the standard normalization for a axial Killing field when the sphere's geometry happens to be axially symmetric.

Owen and collaborators [4] outline such a normalization scheme for the case when u_0^a is constrained to be divergence free on $\mathcal{M} \sim S^2$, and is assumed to vanish only at two isolated points. The orbits of the flow $U_0(t)$ generated by u_0^a are then all circles, each of which closes for some, though not necessarily the same, value of t. The approximate Killing field is normalized such that a certain average of those parameter values t at which each orbit closes is 2π. This reduces to the correct normalization when $u_0^a = \xi^a$ is a genuine axial Killing field, whose orbits all close at $t = 2\pi$ by definition.

Here we propose an alternate normalization for the approximate Killing field on a topological 2-sphere. Our proposal is (a) more local in that it does not require a global average over an entire sphere, (b) does not need to restrict to divergence-free vector fields, and (c) does not need to assume that u_0^a has exactly two isolated zeroes.

Suppose for the moment that $(\mathcal{M} \sim S^2, g_{ab})$ admits an axial Killing field ξ^a. The defining characteristic of an *axial* Killing field is that it vanishes at at least one $p_0 \in \mathcal{M}$. The flow $\Xi(t)$ generated by ξ^a has period $t = 2\pi$ near any such p_0 if and only if

$$\left[\nabla_a \xi^b \cdot \nabla_b \xi^a\right]_{p_0} = -2. \tag{8}$$

This is so because the flow near p_0 generates a corresponding flow in the tangent space

$T_{p_0}\mathcal{M}$, which is a family of SO(2) rotations for the *Euclidean* metric $g_{ab}(p_0)$ thereon. This family of SO(2) rotations closes at $t = 2\pi$ if and only if its matrix generator is properly normalized, which is precisely the condition (8). We show in a companion paper [14] that this convention to normalize ξ^a at *one* fixed point p_0 necessarily selects the desired normalization of the Killing field throughout. Moreover, this result is independent of which p_0 one uses.

It is well known that any vector field, and in particular the approximate Killing field u_0^a for an arbitrary 2-sphere metric g_{ab}, must vanish at at least one $p_0 \in \mathcal{M}$. We therefore can choose the constant scaling of u_0^a throughout $\mathcal{M}$ such that

$$\left[\nabla_a u_0^b \cdot \nabla^a u_{0b}\right]_{p_0} = 2 \tag{9}$$

at that one point. If there is more than one such point, then one may either just pick one, though then the normalization is unlikely to be independent of that choice, or else average the left side of (9) over all such points. (It may be necessary to assume that there are only finitely many zeroes $p_0 \in \mathcal{M}$ of u_0^a, which certainly seems likely for the lowest eigenfields of an elliptic operator. We have no proof of this, however.) Whichever approach one chooses, the normalization of u_0^a given by (9) has both of the desired properties that (a) it exists for any Riemannian 2-sphere geometry, and (b) it gives the correct scaling of $u_0^a = \xi^a$ when that geometry happens to be axially symmetric.

4. Perturbations of the Killing Laplacian and its eigenfields

Next we want to show that the approximate Killing vector eigenvalue problem is *stable*, that is, that small changes in g_{ab} induce only small changes in u_0^a. We approach this problem technically by considering a smooth, one-parameter family $\big(\mathcal{M}, g_{ab}(\lambda)\big)$ of Riemannian geometries on a fixed, compact manifold $\mathcal{M}$ without boundary.

As $g_{ab}(\lambda)$ varies with λ, so will any geometric structure derived from it. The relevant formulae for perturbations of these derived structures are the standard ones from linear perturbation theory in general relativity (see, *e.g.*, [8]). To summarize our conventions, the Levi-Civita connection ∇_a changes according to

$$\delta\big(\nabla_a \omega_b\big) = \dot{\nabla}_{ab}{}^c \, \omega_c := \frac{1}{2} g^{cd} \left(\nabla_d \dot{g}_{ab} - 2\nabla_{(a} \dot{g}_{b)d}\right) \omega_c, \tag{10}$$

where ω_b is an arbitrary fixed (*i.e.*, λ-independent) covector field, while the inverse metric g^{ab} and the volume element ϵ change according to

$$\delta\big(g^{ab}\big) = -\dot{g}^{ab} := -g^{ac} g^{bd} \dot{g}_{cd} \qquad \text{and} \qquad \delta\epsilon = \frac{1}{2} \dot{g} \, \epsilon := \frac{1}{2} g^{ab} \dot{g}_{ab} \, \epsilon, \tag{11}$$

respectively. We denote perturbations using either the operator $\delta := \frac{\partial}{\partial\lambda}\big|_{\lambda=0}$, when necessary, or equivalently using a dot accent. Using these standard results, the perturbation of the Killing Laplacian (6) takes the form

$$\dot{\Delta}_{\text{K}} v^a = \underbrace{\left(2\,\delta_d^{(a}\,\dot{g}^{b)c}\right)}_{X^{abc}{}_d} \nabla_b \nabla_c \, v^d + \underbrace{\left(2\,\dot{\nabla}^b{}_c{}^a + g^{ab}\,\dot{\nabla}_{cd}{}^d - \delta_c^a\,\dot{\nabla}_d{}^{db}\right)}_{Y^{ab}{}_c} \nabla_b \, v^c + \underbrace{\left(2\,\nabla_c \dot{\nabla}_b{}^{(ac)}\right)}_{Z^a{}_b} v^b, \tag{12}$$

where v^d is a fixed vector field. If the vector field is not fixed, such as if it is an eigenfield $u_n^a(\lambda)$ of the Killing Laplacian $\Delta_{\mathrm{K}}(\lambda)$ constructed from the metric $g_{ab}(\lambda)$, then of course there are additional terms coming from the variation of the vector field itself.

The process of relating the perturbation (12) of the Killing Laplacian to the perturbations of its eigenfields is more or less the familiar one from elementary quantum mechanics. Perturbing the eigenvalue problem (6) gives

$$\big(\kappa_n - \Delta_{\mathrm{K}}\big)\dot{u}_n^a = \big(\dot{\Delta}_{\mathrm{K}} - \dot{\kappa}_n\big)u_n^a, \tag{13}$$

which formally determines u_n^a up to its component in the eigenspace of Δ_{K} containing u_n^a itself. One difference from the usual quantum mechanical perturbation theory is that the inner product on the Hilbert space $L^2(\mathcal{M}^\bullet, g)$ varies with λ, too. Imposing the usual normalization condition leads to

$$\|u_n(\lambda)\|^2_{L^2(\mathcal{M}^\bullet, g(\lambda))} = 1 \qquad \rightsquigarrow \qquad 2\,\mathrm{Re}\,\langle u_n, \dot{u}_n\rangle_{L^2(\mathcal{M}^\bullet,g)} = -\langle u_n, \mathcal{A}u_n\rangle_{L^2(\mathcal{M}^\bullet,g)}, \tag{14}$$

where $\mathcal{A}$ denotes the algebraic operator on $L^2(\mathcal{M}^\bullet, g)$ derived from the tensor field

$$\dot{\alpha}^a{}_b := \dot{g}^a{}_b + \frac{1}{2}\,\dot{g}\,\delta^a_b. \tag{15}$$

This accounts for the changes of the point-wise inner product in the tangent space and the volume element in the Hilbert-space inner product (7). Collecting these results, and keeping u_n^a real-valued, gives the formal expression

$$\dot{u}_n^a = \left(P_n^\perp \frac{1}{\kappa_n - \Delta_{\mathrm{K}}} P_n^\perp\right)\dot{\Delta}_{\mathrm{K}}\,u_n^a - \frac{1}{2}\,\dot{\alpha}^a{}_b\,u_n^b \tag{16}$$

for the perturbation of the n^{th} eigenfield of the Killing Laplacian.

The operator in braces in (16) is the pseudoinverse of the operator on the left side of (13). That is, it represents convolution with the Green function for the restriction of that operator to the subspace of $L^2(\mathcal{M}^\bullet, g)$ orthogonal to the eigenspace of Δ_{K} with eigenvalue κ_n. We use $P_n^\perp$ to denote the projection onto that orthogonal subspace. The pseudoinverse operator is bounded (in fact, compact), and its operator norm obeys

$$\left\| P_n^\perp \frac{1}{\kappa_n - \Delta_{\mathrm{K}}} P_n^\perp \right\| \leq \frac{1}{\gamma_n}, \qquad \text{where} \qquad \gamma_n := \min(\kappa_n - \kappa_{n-1}, \kappa_{n+1} - \kappa_n) \tag{17}$$

is the gap around κ_n in the spectrum of Δ_{K}. (In the case $n = 0$ of primary interest, its definition is simply $\gamma_0 := \kappa_1 - \kappa_0$.) We are relying on general results for elliptic operators on compact manifolds [13, 15] to show that that spectrum is discrete, and therefore that this bound on the operator norm of the pseudoinverse is meaningful. The important point, however, is that bounding the perturbation $\dot{u}_n^a$ of the eigenfield in Hilbert space amounts to bounding the action of the perturbed Killing Laplacian $\dot{\Delta}_{\mathrm{K}}$ on the unperturbed u_n^a.

5. Bounding the perturbed eigenfields of the Killing Laplacian

This section shows that the action $\dot{\Delta}_{\mathrm{K}} v^a$ of the perturbed Killing Laplacian on an arbitrary vector field is necessarily bounded as long as both $\Delta_{\mathrm{K}} v^a$ and v^a itself have finite L^2 norm. (This will be the case, for example, if v^a lies in the domain of Δ_{K} in Hilbert space.) This implies further that the norm of the perturbation $\dot{u}^a_n$ of any eigenfield of the Killing Laplacian is less than a certain constant C_n times the norm of the unperturbed u^a_n. Moreover, this C_n scales in direct proportion to the metric perturbation $\dot{g}_{ab}$. For any given n, we therefore can make (the norm of) $\dot{u}^a_n$ as small as we like by taking $\dot{g}_{ab}$ sufficiently small. In other words, each eigenfield of the Killing Laplacian depends *continuously*, in the natural L^2 topology on Hilbert space, on the Riemannian metric g_{ab}. This gives the precise mathematical sense in which the approximate Killing field determined by the eigenvalue problem (6) is stable.

The following proof of these claims is organized in three steps. The first is just to use the triangle inequality and certain point-wise bounds on the tensor fields $X^{ab}{}_c{}^d$, $Y^b{}_c{}^d$ and $Z_c{}^d$ from (12) to show that

$$\left\|\dot{\Delta}_{\mathrm{K}} v\right\|_{L^2(\mathcal{M}^\bullet)} \leq C_X \left\|\nabla\nabla v\right\|_{L^2(\mathcal{M}_{\bullet\bullet}{}^\bullet)} + C_Y \left\|\nabla v\right\|_{L^2(\mathcal{M}_\bullet{}^\bullet)} + C_Z \left\|v\right\|_{L^2(\mathcal{M}^\bullet)}. \tag{18}$$

The new L^2 norms appearing here are just the natural ones for tensor fields with the given index structure under the *background* metric $g_{ab}(0)$, as in (7). (All norms and inner products from this point forward derive from the background metric, which accordingly we no longer denote explicitly.) The constants C_X, C_Y and C_Z depend on the tensor fields from (12), and scale in direct proportion to the metric perturbation $\dot{g}_{ab}$. We derive (18) in subsection 5.1, including specific values for C_X, C_Y and C_Z.

The most divergent term on the right side of (18) typically will be the first, which involves the (covariant) Hessian $\nabla_a \nabla_b v^c$ of the vector field. The second step of our calculation is therefore to bound this Hessian such that

$$\begin{aligned}\left\|\nabla\nabla v\right\|^2_{L^2(\mathcal{M}_{\bullet\bullet}{}^\bullet)} \leq \left\|\Delta_{\mathrm{K}} v\right\|^2_{L^2(\mathcal{M}^\bullet)} + \left\|\mathrm{Riem}\, v\right\|^2_{L^2(\mathcal{M}_{\bullet\bullet}{}^\bullet)} \\ + C_U \left\|\nabla v\right\|^2_{L^2(\mathcal{M}_\bullet{}^\bullet)} + C_V \left\|v\right\|^2_{L^2(\mathcal{M}^\bullet)}.\end{aligned} \tag{19}$$

The second term on the right here denotes the norm of the tensor field $v^a R_{abc}{}^d$, which actually is equal to the Hessian $\nabla_b \nabla_c \xi^d$ if $v^a = \xi^a$ happens to be a Killing field. This provides a nice check on this inequality since the first term on the right vanishes in that case, and one has only some very general bounds on the constants C_U and C_V in the second line. Indeed, those constants arise in much the same way as the corresponding constants in (18), though they depend only on the background geometry g_{ab} and not its perturbation $\dot{g}_{ab}$. We derive (19) in subsection 5.2.

The first two steps of the calculation show that the action of the perturbed Killing Laplacian $\dot{\Delta}_{\mathrm{K}}$ on an arbitrary vector field v^a has bounded L^2 norm if we can control (a) the L^2 norm of $\Delta_{\mathrm{K}} v^a$, and (b) the L^2 norm of the first-order derivatives $\nabla_a v^b$. The third step of the calculation shows that the first-order derivatives can also be bounded

in terms of $\Delta_{\mathrm{K}} v^a$ according to

$$\|\nabla v\|^2_{L^2(\mathcal{M}_{\bullet}{}^{\bullet},g)} \leq \langle v, \Delta_{\mathrm{K}} v\rangle_{L^2(\mathcal{M}^{\bullet},g)} + C_{\mathrm{Ric}}\, \|v\|^2_{L^2(\mathcal{M}^{\bullet},g)}, \tag{20}$$

where C_{Ric} bounds the linear operator $v^a \mapsto R^a{}_b\, v^b$ on Hilbert space derived from the Ricci tensor of the background metric. Again, it is independent of the perturbation $\dot{g}_{ab}$. We establish this last result in subsection 5.3.

At this stage we (will) have shown that, although the perturbed Killing Laplacian $\dot{\Delta}_{\mathrm{K}}$ is certainly an unbounded operator on Hilbert space, the "unbounded part" of its action on any given vector field v^a is *relatively* bounded by the action $\Delta_{\mathrm{K}} v^a$ of the Killing Laplacian itself. (One could insert some technical language at this point regarding the domains of these operators, but in fact this argument shows that it is natural simply to choose the domain of $\dot{\Delta}_{\mathrm{K}}$ to coincide with that of Δ_{K}.) The action $\Delta_{\mathrm{K}} u^a_n = \kappa_n\, u^a_n$ of the Killing Laplacian on any of its own eigenfields is certainly bounded. We therefore conclude from (16) that the perturbation $\dot{u}^a_n$ of any eigenfield of Δ_{K} satisfies

$$\|\dot{u}_n\|_{L^2(\mathcal{M}^{\bullet})} \leq C_n\, \|u_n\|_{L^2(\mathcal{M}^{\bullet})} \tag{21}$$

for a constant C_n depending on the unperturbed eigenvalue κ_n, the unperturbed spectral gap γ_n around that eigenvalue, and a number of constants C like those described above depending on g_{ab} and $\dot{g}_{ab}$, but not on n. We establish this result in detail in subsection 5.4, where we give a specific formula for the C_n.

5.1. The Hessian and lower-order terms bound the perturbed Killing Laplacian

Applying the triangle inequality to the perturbed Killing Laplacian (12) gives

$$\begin{aligned}\left\|\dot{\Delta}_{\mathrm{K}} v\right\|_{L^2(\mathcal{M}^{\bullet})} &:= \|\mathcal{X}\nabla\nabla v + \mathcal{Y}\nabla v + \mathcal{Z} v\|_{L^2(\mathcal{M}^{\bullet})} \\ &\leq \|\mathcal{X}\nabla\nabla v\|_{L^2(\mathcal{M}^{\bullet})} + \|\mathcal{Y}\nabla v\|_{L^2(\mathcal{M}^{\bullet})} + \|\mathcal{Z} v\|_{L^2(\mathcal{M}^{\bullet})},\end{aligned} \tag{22}$$

where $\mathcal{X}$, $\mathcal{Y}$ and $\mathcal{Z}$ denote the obvious algebraic operators on Hilbert space derived from the tensor fields X, Y and Z in (12), respectively.

The next step is more or less the same for each of the three terms in (22). Let's consider the first in detail, which is

$$\|\mathcal{X}\nabla\nabla v\|^2_{L^2(\mathcal{M}^{\bullet})} = \left\langle \nabla\nabla v, \mathcal{X}^\dagger\mathcal{X}\nabla\nabla v\right\rangle_{L^2(\mathcal{M}_{\bullet\bullet}{}^{\bullet})}. \tag{23}$$

The operator $\mathcal{X}^\dagger\mathcal{X}$ is obviously non-negative. It is also bounded because $\mathcal{M}$ is compact by assumption and both $\mathcal{X}$ and $\mathcal{X}^\dagger$ act algebraically. To see this, note that the tensor field underlying $\mathcal{X}^\dagger$ is just the point-wise adjoint of that underlying $\mathcal{X}$ itself:

$$X^\dagger{}_{ab}{}^c{}_d\, w^d := g_{aj}\, g_{bk}\, g^{cl}\, X^{ijk}{}_l\, g_{id}\, w^d \qquad \text{or} \qquad X^{\dagger A}{}_d := g^{AB}\, g_{cd}\, X^c{}_B. \tag{24}$$

We use uppercase indices in the latter expression to denote "vectors" in the appropriate tensor product of tangent and co-tangent spaces at each point of $\mathcal{M}$. The product

$X^{\dagger A}{}_b\, X^b{}_C$ at each $p \in \mathcal{M}$ is then a non-negative, Hermitian map from that tensor product space to itself. In fact, each of these maps must have a substantial kernel because $X^i{}_J$ itself maps each eigenvector of $(X^\dagger X)^A{}_C$ to an eigenvector of $(XX^\dagger)^a{}_c$ having the same eigenvalue, and conversely $X^{\dagger I}{}_j$ maps eigenvectors in the opposite direction. It follows that each $(X^\dagger X)^A{}_C$ has at most as many non-zero eigenvalues as the dimension of $\mathcal{M}$, and each of these is also an eigenvalue of $(XX^\dagger)^a{}_c$ at the same point. Define $M(p)$ to be the largest such eigenvalue at each $p \in \mathcal{M}$. This function must be smooth since the tensor fields used to define it are smooth. It therefore attains a finite maximum value on the compact manifold $\mathcal{M}$. This "maximum maximum" eigenvalue bounds the Hilbert space operator $\mathcal{X}^\dagger\mathcal{X}$ such that

$$\|\mathcal{X}\nabla\nabla v\|_{L^2(\mathcal{M}^\bullet)} \leq C_X \, \|\nabla\nabla v\|_{L^2(\mathcal{M}_{\bullet\bullet}{}^\bullet)}. \tag{25}$$

We have shown that the bounding constant C_X here is given by

$$\begin{aligned} C_X^2 := \max_{p\in\mathcal{M}} M(p) &:= \max_{p\in\mathcal{M}} \max_{s\in S_p\mathcal{M}} \Big[s_a \, X^{abc}{}_d \, X_{ebc}{}^d \, s^e \Big](p) \\ &\leq \max_{p\in\mathcal{M}} \Big[X^{abc}{}_d \, X_{abc}{}^d \Big](p), \end{aligned} \tag{26}$$

where we have used the metric to raise and lower indices and the inner maximum on the first line is over the unit sphere $S_p\mathcal{M}$ in each tangent space $T_p\mathcal{M}$. The final inequality arises because the largest eigenvalue at each point is less than the sum of all the (non-negative) eigenvalues, which of course is just the trace of the operator. This trace might be simpler to calculate and maximize over $\mathcal{M}$ in practice.

We can bound the other two terms on the right side of (22) in a completely analogous manner. The overall bound (18) then follows directly. Note that the tensor fields $X^{abc}{}_d$, $Y^{ab}{}_c$ and $Z^a{}_b$ from (12) are all linear in the perturbation $\dot{g}_{ab}$ of the metric. It follows that the bounding constants C_X, C_Y and C_Z in (18) all scale homogeneously under a reparameterization of the family of geometries $g_{ab}(\lambda)$ underlying the perturbation. The left side of (18) scales in the same manner, of course, so the result holds regardless of how the parameter λ is defined geometrically.

5.2. The Killing Laplacian and lower-order terms bound the Hessian

We begin by establishing a couple local differential identities involving the Killing Laplacian and related operators. The first is a Weitzenböck identity relating Δ_{K} to the *Laplace–Beltrami operator* $\Delta_{\mathrm{B}} := -\nabla_a\nabla^a$ acting on vector fields:

$$\begin{aligned} \Delta_{\mathrm{K}} v^a := -\nabla_b\nabla^b \, v^a - \nabla_b\nabla^a \, v^b &= \Delta_{\mathrm{B}} v^a - \nabla^a\nabla_b \, v^b - R^a{}_b \, v^b \\ &= (\Delta_{\mathrm{B}} v - \operatorname{grad}\operatorname{div} v - \operatorname{Ric} v)^a. \end{aligned} \tag{27}$$

The index-free notation on the second line will be useful for the manipulations in Hilbert space that follow. The second set of results we will need is

$$0 = 2\,\nabla_a\nabla_b\nabla^{[a} v^{b]} = \operatorname{div}(\mathrm{d}^\dagger \mathrm{d} v) = \operatorname{div}(\Delta_{\mathrm{B}} v + \operatorname{grad}\operatorname{div} v + \operatorname{Ric} v), \tag{28}$$

where we use the metric to raise and lower indices as needed, d in the third expression denotes the exterior derivative, and $\mathrm{d}^\dagger$ denotes its adjoint in the natural Hilbert inner product on the space of 2-form fields. Note that the final equality also holds without the initial divergence, as a relation between vector fields.

We now compute the L^2 norm of $\Delta_{\mathrm{K}} v^a$ as follows. First, substitute for Δ_{K} using (27) and expand. Second, use integrations by parts and the identities (28) to simplify the resulting cross-terms, leaving

$$\begin{aligned}\|\Delta_{\mathrm{K}} v\|^2_{L^2(\mathcal{M}^\bullet,g)} = \|\Delta_{\mathrm{B}} v\|^2_{L^2(\mathcal{M}^\bullet)} - 2\,\mathrm{Re}\,\langle \mathrm{d}v, \mathrm{d}\,\mathrm{Ric}\,v\rangle_{L^2(\mathcal{M}_{[\bullet\bullet]})} \\ + 3\,\|\mathrm{grad}\,\mathrm{div}\,v + \mathrm{Ric}\,v\|^2_{L^2(\mathcal{M}^\bullet)}. \end{aligned} \tag{29}$$

Note that the second term on the right here involves at most first-order derivatives of v^a. It can also be written in the form

$$\begin{aligned}\langle \mathrm{d}v, \mathrm{d}\,\mathrm{Ric}\,v\rangle_{L^2(\mathcal{M}_{[\bullet\bullet]})} &:= \frac{1}{2}\oint_{\mathcal{M}} \big(2\,\nabla^{[a}\,\bar{v}^{b]}\big)\big(2\,\nabla_{[a}R_{b]c}\cdot v^c + 2\,R_{d[b}\,g_{a]c}\,\nabla^c\,v^d\big)\,\epsilon \\ &= \oint_{\mathcal{M}} \big(2\,g^{c[a}\,R^{b]d}\big)\big(\nabla_a\,\bar{v}_b\big)\big(\nabla_c\,v_d\big)\,\epsilon - \oint_{\mathcal{M}} \big(\nabla_d\,R_{abc}{}^d\big)\big(\nabla^a\,\bar{v}^b\big)\,v^c\,\epsilon,\end{aligned} \tag{30}$$

where we have used the partially contracted Bianchi identity $2\,\nabla_{[a}\,R_{b]c} = -\nabla_d\,R_{abc}{}^d$ in the second line. Third, the squared norm of $\Delta_{\mathrm{B}} v^a$ is related to the squared norm of the Hessian of v^a by additional lower-order derivative terms:

$$\begin{aligned}\|\Delta_{\mathrm{B}} v\|^2_{L^2(\mathcal{M}^\bullet)} = \|\nabla\nabla v\|^2_{L^2(\mathcal{M}_{\bullet\bullet}{}^\bullet)} + \oint_{\mathcal{M}} \big(R^{ac}\,g^{bd}\big)\big(\nabla_a\,\bar{v}_b\big)\big(\nabla_c\,v_d\big)\,\epsilon \\ + \oint_{\mathcal{M}} \big(\nabla_d\,R_{bca}{}^d\big)\big(\nabla^a\,\bar{v}^b\big)\,v^c\,\epsilon - \big\|\mathrm{Riem}\,v\big\|^2_{L^2(\mathcal{M}_{\bullet\bullet}{}^\bullet)}.\end{aligned} \tag{31}$$

The last term here is the squared L^2 norm of $v^a\,R_{abc}{}^d$. As mentioned above, this is equal to the Hessian $\nabla_b\nabla_c\,\xi^d$ when $v^a = \xi^a$ is a Killing field. Finally, insert (31) and (30) into (29) and collect terms to find

$$\begin{aligned}\|\nabla\nabla v\|^2_{L^2(\mathcal{M}_{\bullet\bullet}{}^\bullet)} &\le \|\nabla\nabla v\|^2_{L^2(\mathcal{M}_{\bullet\bullet}{}^\bullet)} + 3\,\|\mathrm{grad}\,\mathrm{div}\,v + \mathrm{Ric}\,v\|^2_{L^2(\mathcal{M}^\bullet)} \\ &= \|\Delta_{\mathrm{K}} v\|^2_{L^2(\mathcal{M}^\bullet)} + \|\mathrm{Riem}\,v\|^2_{L^2(\mathcal{M}_{\bullet\bullet}{}^\bullet)} - \oint_{\mathcal{M}} \underbrace{\big(\nabla_a\nabla_c\,R^{abcd}\big)}_{V^{bd}}\,\bar{v}_b\,v_d\,\epsilon \\ &\quad + \oint_{\mathcal{M}} \overbrace{\big(2\,g^{ac}\,R^{bd} - R^{ac}\,g^{bd} - g^{ad}\,R^{cb} - g^{cb}\,R^{ad}\big)}^{U^{abcd}}\big(\nabla_a\,\bar{v}_b\big)\big(\nabla_c\,v_d\big)\,\epsilon \\ &\le \|\Delta_{\mathrm{K}} v\|^2_{L^2(\mathcal{M}^\bullet)} + \|\mathrm{Riem}\,v\|^2_{L^2(\mathcal{M}_{\bullet\bullet}{}^\bullet)} \\ &\quad + \|\nabla v\|_{L^2(\mathcal{M}_\bullet{}^\bullet)}\,\|\mathcal{U}\nabla v\|_{L^2(\mathcal{M}_\bullet{}^\bullet)} + \|v\|_{L^2(\mathcal{M}^\bullet,g)}\,\|\mathcal{V}v\|_{L^2(\mathcal{M}^\bullet)}.\end{aligned} \tag{32}$$

The last line of this calculation uses the Schwarz inequality. The operators $\mathcal{U}$ and $\mathcal{V}$ on Hilbert space are defined in the obvious way from the tensor fields $U_a{}^{bc}{}_d$ and $V^b{}_d$, respectively. These algebraic operators are bounded, and their bounds can be computed by the same argument that led to (25). Our main result (19) follows. Unlike (18), however, the bounding constants C_U and C_V in this case depend only on the background geometry g_{ab}.

5.3. The Killing Laplacian and algebraic terms bound the first-order derivatives

Combining our previous results (18) and (19) bounds the L^2 norm of $\dot{\Delta}_{\mathrm{K}} v^a$ in terms of the L^2 norms of $\Delta_{\mathrm{K}} v^a$, $\nabla_a v^b$ and v^a itself. We now show that $\|\nabla v\|_{L^2(\mathcal{M}_\bullet{}^\bullet)}$ can also be bounded in terms of $\|\Delta_{\mathrm{K}} v\|_{L^2(\mathcal{M}^\bullet)}$. The idea is just to use the identity (27) and integration by parts to show that

$$\langle v, \Delta_{\mathrm{K}} v\rangle_{L^2(\mathcal{M}^\bullet)} = \|\nabla v\|^2_{L^2(\mathcal{M}_\bullet{}^\bullet)} + \|\mathrm{div}\, v\|^2_{L^2(\mathcal{M})} + \langle v, \mathrm{Ric}\, v\rangle_{L^2(\mathcal{M}^\bullet)}. \tag{33}$$

Throwing away the norm of the divergence and using the Schwarz inequality for the term involving the Ricci tensor gives the inequality we seek,

$$\|\nabla v\|^2_{L^2(\mathcal{M}_\bullet{}^\bullet)} \leq \langle v, \Delta_{\mathrm{K}} v\rangle_{L^2(\mathcal{M}^\bullet)} + \|v\|_{L^2(\mathcal{M}^\bullet)} \|\mathrm{Ric}\, v\|_{L^2(\mathcal{M}^\bullet)}. \tag{34}$$

The expectation value on the right is non-negative because the Killing Laplacian is non-negative. Furthermore, the term involving the Ricci tensor can be bounded via the argument that led to (25). The result (20) follows. Thus, the action of the perturbed Killing Laplacian on an arbitrary vector field v^a can be bounded in terms of the action of the Killing Laplacian itself and terms algebraic in v^a.

5.4. Dénouement

Nesting (19) and (20) inside (18) gives the bound

$$\begin{aligned}\left\|\dot{\Delta}_{\mathrm{K}} v\right\|_{L^2(\mathcal{M}^\bullet)} \leq C_X \Big(&\|\Delta_{\mathrm{K}} v\|^2_{L^2(\mathcal{M}^\bullet)} + C_U \,\langle v, \Delta_{\mathrm{K}} v\rangle_{L^2(\mathcal{M}^\bullet, g)} \\ &+ \left(C_V + C_U\, C_{\mathrm{Ric}} + C^2_{\mathrm{Riem}}\right) \|v\|^2_{L^2(\mathcal{M}^\bullet)} \Big)^{1/2} \\ &+ C_Y \left(\langle v, \Delta_{\mathrm{K}} v\rangle_{L^2(\mathcal{M}^\bullet, g)} + C_{\mathrm{Ric}} \,\|v\|^2_{L^2(\mathcal{M}^\bullet, g)} \right)^{1/2} + C_Z \,\|v\|_{L^2(\mathcal{M}^\bullet)}.\end{aligned} \tag{35}$$

on the L^2 norm of $\dot{\Delta}_{\mathrm{K}} v^a$ for an arbitrary vector field v^a. This result is of some interest in its own right, but our main objective now is to take $v^a = u^a_n$ to be an eigenfield of Δ_{K}. We then can combine the formal expression (16) for the perturbed eigenfield $\dot{u}^a_n$ with the operator norm (17) of the pseudoinverse to yield (21) with

$$C_n := \frac{C_X \sqrt{\kappa_n^2 + C_U\, \kappa_n + C_V + C_U\, C_{\mathrm{Ric}} + C^2_{\mathrm{Riem}}}}{\gamma_n} + \frac{C_Y \sqrt{\kappa_n + C_{\mathrm{Ric}}}}{\gamma_n} + \frac{C_Z}{\gamma_n} + \frac{C_A}{2}. \tag{36}$$

The individual constants C here are all of the form (26) for the tensor fields $X^{a\cdot bc}{}_d$, $Y^{a\cdot b}{}_c$ and $Z^a{}_b$ from (12); $U_a{}^{b\cdot c}{}_d$ and $V^a{}_b$ from (32); the curvatures $R^a{}_b$ and $R_{a\cdot bc}{}^d$; and lastly $\dot{\alpha}^a{}_b$ from (15). We have used a dot to separate indices, where necessary, to show how each of these tensor fields should be interpreted as a mapping between tensor products of tangent and cotangent spaces. The grouping of indices is important for calculating the "maximum maximum" eigenvalue of each tensor field as on the first line (26), but not for merely calculating the trace on the second line.

We conclude this detailed discussion with three comments.

First, the constants C_n from (36) are homogeneous of degree one in the metric perturbation $\dot{g}_{ab}$ (and its derivatives). Note that $X^{abc}{}_d$, $Y^{ab}{}_c$, and $Z^a{}_b$ from (12), as well as $\dot{\alpha}^a{}_b$ from (15), scale proportionately under a (constant) rescaling of $\dot{g}_{ab}$. The other "C" constants appearing in (36) derive solely from the background geometry g_{ab}.

Second, the constants C_n from (36) vary with n, and increase like κ_n/γ_n for large n. One expects the eigenvalue gap γ_n to increase with n, but not as rapidly as κ_n itself. Thus, for example, the fractional perturbations in the eigenvalues themselves,

$$\frac{\dot{\kappa}_n}{\kappa_n} = \frac{\langle u_n, \dot{\Delta}_{\mathrm{K}} u_n \rangle}{\kappa_n \left\| u_n \right\|^2} \leq \frac{C_n}{\kappa_n}, \tag{37}$$

vanish for asymptotically large n. Nonetheless, the C_n themselves diverge in that limit. Therefore, while each eigenfield u^a_n varies continuously with g_{ab} in the sense of (21), the *set* of eigenfields is not *uniformly* continuous in g_{ab}. This may raise some technical difficulties in certain applications, but our objective here is just to bound the variations only of the lowest eigenfield u^a_0. Thus, what matters here is just the one constant C_0, and the subtlety regarding uniform continuity is irrelevant.

Third, we have shown that u^a_0 varies continuously with g_{ab} in a loose sense, but have not stated clearly the topology on the space of metrics for which this continuity holds strictly. Though this is a somewhat subtle issue, it certainly is clear that the relevant topology does not derive from a Hilbert-type (L^2) norm. Rather, the constants C_X, C_Y, C_Z, and C_A in (36) arise by maximizing quantities over all of $\mathcal{M}$ as in (26). The appropriate topology on the space of metrics should therefore be based on a supremum, or L^∞-type, norm of $\dot{g}_{ab}$. It should also have a Sobolev-type character, depending not only on $\dot{g}_{ab}$ but also on its first two derivatives. The natural candidate is

$$\|\dot{g}\|^2_{W^\infty_2(\mathcal{M}_{(\bullet\bullet)},g)} := \max_{p\in\mathcal{M}} \Big[\dot{g}^{ab}\,\dot{g}_{ab} + \ell^2 \big(\nabla_a \dot{g}_{bc}\big)\big(\nabla^a \dot{g}^{bc}\big) + \ell^4 \big(\nabla_{(a}\nabla_{b)}\dot{g}_{cd}\big)\big(\nabla^{(a}\nabla^{b)}\dot{g}^{cd}\big) \Big], \tag{38}$$

where ℓ is a length scale that may depend on g_{ab}. (For example, it might derive from the total volume of $\mathcal{M}$ in the background geometry.) Conversely, we have established L^2-type bounds on the perturbations $\dot{u}^a_n$, but not L^∞-type (*i.e.*, supremum) bounds. The latter are irrelevant, however, for calculating perturbations in quasi-local invariants like (2) or (3).

6. Conclusions

This paper considers a family of smooth metrics $g_{ab}(\lambda)$ depending smoothly on a real parameter λ. Each metric defines a Killing Laplacian $\Delta_{\mathrm{K}}(\lambda)$, which admits a complete basis of eigenfields $u_n(\lambda)$ in the Hilbert space $L^2(\mathcal{M}^\bullet, g)$ of vector fields on a compact manifold $\mathcal{M}$. (The *topology* of this Hilbert space, though of course not its inner product, is independent of the smooth metric $\mathcal{M}$ on the compact manifold $\mathcal{M}$.) We have shown that each eigenfield $u^a_n(\lambda)$ varies continuously with g_{ab} in the L^2 sense of (21), where the size of a metric perturbation is measured in the W^∞_2 sense of (38). In particular, the

lowest-order eigenfield u_0^a, which is a natural candidate to define an approximate Killing field on a manifold with no genuine symmetries, depends continuously on the metric.

The natural application of this result is to the problem of measuring the angular momentum of a black hole in general relativity. The manifold $\mathcal{M}$ is a topological 2-sphere in this case, and the metric g_{ab} on it is arbitrary. The result is useful because in practical applications, say in numerical relativity, the actual horizon may be difficult or expensive to locate exactly. Standard horizon-finding schemes [11] help to locate the horizon approximately, and the continuous dependence of the approximate Killing eigenfield u_0^a on the metric implies that the corresponding quasi-local angular momentum of the approximate horizon will approximate that of the actual horizon. This conclusion is not unexpected, of course, but it is useful to establish a solid theoretical basis for the claim. Similarly, the general continuity result we have established can be used to bound the growth of the angular momentum of a black hole under dynamic evolution. This could be useful, for example, in the study of absorption processes involving dynamical horizons [10] at or near extremality. Finally, in each of these applications we have proposed in (9) a novel, and purely local means of normalizing the approximate Killing field on a topological 2-sphere.

Acknowledgements

The authors thank their colleagues at Florida Atlantic University for stimulating dicusssions of this work. SW also gratefully acknowledges support under National Science Foundation Grant DGE:0638662 while part of this work was completed.

Bibliography

[1] O. Dreyer, B. Krishnan, E. Schnetter and D. Shoemaker. Introduction to isolated horizons in numerical relativity. *Phys. Rev. D* **67** 024018 (2003).

[2] A.I. Harte. Approximate spacetime symmetries and conservation laws. *Class. Quantum Grav.* **25** 205008 (2008).

[3] G.B. Cook and B.F. Whiting. Approximate Killing vectors on S^2. *Phys. Rev. D* **76** 041501(R) (2007).

[4] G. Lovelace, R. Owen, H.P. Pfeiffer and T. Chu. Binary-black-hole initial data with nearly extremal spins. *Phys. Rev. D* **78** 084017 (2008).

[5] C. Beetle. Approximate Killing Fields as an Eigenvalue Problem. Preprint `arXiv:0808.1745 [gr-qc]`.

[6] R.A. Matzner. Almost Symmetric Spaces and Gravitational Radiation. *J. Math. Phys.* **9** 1657–1668 (1968).

[7] A. Komar. Covariant Conservation Laws in General Relativity. *Phys. Rev.* **113** 934–936 (1959).

[8] R.M. Wald. *General Relativity*. University of Chicago Press, Chicago, 1984.

[9] A. Ashtekar and B. Krishnan. Dynamical horizons and their properties. *Phys. Rev. D* **68** 104030 (2003).

[10] A. Ashtekar and B. Krishnan. Isolated and Dynamical Horizons and Their Applications. *Living Rev. Relativity* **7** 10 (2004).
`http://www.livingreviews.org/lrr-2004-10` (cited 30 December 2013).

[11] J. Thornburg. Event and Apparent Horizon Finders for $3 + 1$ Numerical Relativity. *Living Rev.*

Relativity **10** 3 (2007).
`http://www.livingreviews.org/lrr-2007-3` (cited 30 December 2013).

[12] K. Yano and S. Bochner. *Curvature and Betti Numbers*. Princeton University Press, Princeton, 1953.

[13] W. McLean. *Strongly elliptic systems and boundary integral equations*. Cambridge University Press, Cambridge, 2000.

[14] C. Beetle and S. Wilder. A note on axial symmetries. Preprint `arXiv:?? [gr-qc]`.

[15] Y. Choquet-Bruhat. *General relativity and the Einstein equations*. Oxford University Press, Oxford, 2009.